Complex Adaptive Systems: Energy, Information and Intelligence: Papers from the 2011 AAAI Fall Symposium (FS-11-03)

A Complex Adaptive Systems Investigation of the Social-Ecological Dynamics of Three Fisheries

**P. S. Hayes, J. Wilson, C. B. Congdon, L Yan, J. Hill,
J. Acheson, Y. Chen, C. Cleaver, A. Hayden, T. Johnson,
M. Kersula, G. Morehead, R. Steneck**

School of Marine Science, University of Maine, Orono Maine
Hayes, P.S (peter hayes@maine.edu); Wilson,J. (jwilson@maine.edu); Yan, L.(Liying_Yan@umit maine.edu);
Hill, J.(jack_hill@umit maine.edu); Acheson, J.(Acheson@maine.edu); Chen, Y (ychen@maine.edu);
Cleaver, C.(caitlin.cleaver@maine.edu); Hayden, A.(acjhayden@comcast.net); Johnson, T. (Teresa.johnson@maine.edu);
Kersula, M.(Michael kersula@maine.edu); Morehead, G.(graham.morehead@umit maine.edu); Steneck, R. (steneck@maine.edu)

Department of Computer Science, University of Southern Maine
Congdon, C.B. (congdon@usm maine.edu)

Abstract

In this paper we describe a complex adaptive systems model of interactions between coupled human and natural system. We use learning classifier systems to create adaptive agents in a simulation of the Maine lobster fishery to explore the relationships among ecological, economic, and social characteristics. Our hypothesis is that the cost of information and learning drives agents' decisions to compete or co-operate and, consequently, the emergence of long-term relationships. Initial results provide tentative support for the hypothesis and the ability of this model to provide insight into the dynamics of individual interactions and the social relationships that emerge from those interactions.

Introduction

In a complex, competitive society, a self-interested individual with limited resources must continually make decisions regarding the allocation of those resources. This decision process is constrained by the time and effort dedicated to the acquisition of the knowledge that allows the individual to exploit the surrounding natural and artificial environment in pursuit of personal goals. The vast array of available knowledge and the equally large variation in the value of that knowledge on an individual and temporal basis creates an extremely complex problem for individuals. The knowledge gained (or ignored or lost) exerts influences over the subsequent actions, behavior, and success or failure of individuals; that is, as agents learn about the results of their actions, their new knowledge modifies the value of knowledge previously gained.

Our social hypothesis is that this pursuit of costly knowledge by a large number of self-interested agents creates regular relationships that constrain individual behavior and lead to the emergence of self-organized social structures and dynamics, and that insight into these relationships and the factors affecting them can be critical in better understanding the design and potential outcome of policy intended to regulate behavior. Given the complexity of influences, the lack of control, and the difficulty of studying real human agents in real-life situations, we are developing models by which to investigate aspects of this process using adaptive computer agents. The following sections describe the philosophical background, the ongoing investigations, and some future directions to be pursued while investigating the adaptation of agents to complex human and natural environments from both individual and social perspectives.

Economic Incentives for an Investigation of Behavior in Resource Exploitation

Economists, including Stiglitz (2000), Arthur (1997), and Hayek (1945), speak of economies and, particularly, markets as responding to information gained from a myriad of sources. Although the mechanism of this response is not usually explicit, it is often assumed to involve focused and conscious learning and adaptation by human participants whose individual actions combine to produce the aggregate market behavior seen. Hayek (1945) possibly comes closest to describing a general mechanism by which markets (and the individuals comprising them) can incorporate the huge and varied information sources available. His description involves a hierarchy of nearly decomposable nested spheres of individual knowledge. Individuals economize on their knowledge acquisition by maintaining limited spheres within which their knowledge level is high and they are well-informed. Outside of these spheres of knowledge, individuals are not well-informed. However, the prices and quantities of goods and services entering and leaving their personal spheres provide the necessary information by which to act in their own best interests and to coordinate their activities with others. The result enables the creation and functioning of even the large and global markets existing today.

However appealing at an intuitive level, Hayek's description is lacking in details, and provides little insight into the individual adaptation leading to the markets' self-organized social and economic activities. This 'low-level' understanding is becoming of increasing importance in resource management decisions as we come to question the sustainability of our use of natural resources and realize the significance of fine-scale individual and small group behavior to the success of policy. Ostrom (1990, 2009) has done path-breaking work on the social and bio-physical conditions that facilitate successful collective action, that is, mutual restraint that sustains the resource or activity in question. We are interested in the circumstances that create those conditions. In that light we address three important questions:

- How do individuals with limited time and resources learn and adapt selectively to a large, complex, and competitive environment?
- How do individuals learn and develop restrained behavior necessary for the emergence of social structure and dynamics?
- What conditions lead to emerging structures and dynamics that might contribute to self-governing resource use or the success of management policies?

Gulf of Maine Fisheries as Exemplars of Human Decision-making

In a unique effort, our broadly interdisciplinary group has been developing adaptive agent based models (A-ABMs) of simple human societies exploiting a complex environment. Our A-ABM differs from regular ABM in that each agent incorporates a form of machine learning and artificial intelligence – a modified Learning Classifier System (LCS) – to imbue the agents with individual learning and adaptation capabilities. This allows agents to learn and to adapt their behavior in response to interactions with both the environment and with other agents. This evolutionary approach changes the domain specific knowledge required of the modeler; instead of requiring knowledge of the agents conditional adaptive behavior as in conventional models, i.e., the specification of an equation or behavioral rules, the modeler must acquire knowledge about the kind of information, the actions and the feedback real world agents use. The evolutionary computation then searches out well adapted rules.

The settings for the models being developed consist of three local fishing industries from the Gulf of Maine: the lobster fishery, the sea urchin fishery, and the groundfish fishery. A model of the lobster fishery has already been developed and is being investigated, while models of the two additional fisheries (and a framework in which to better implement future models) are being developed. The three fisheries create very different learning problems for fishermen and lead to very different social structure.

Facets of these three learning problems have traits common to all complex human systems, including:

- long-term regularities, particularly at large scales, in the distribution of a resource, created largely by ecological factors;
- fine-scale irregularities and variability in distribution, in response to ecological variability and the actions of the human fishers themselves;
- the opportunity to both compete and cooperate to achieve goals;

Additionally, agents are provided with the challenges of adapting to

- a complex knowledge environment in which information may be homogeneous or heterogeneous, ephemeral or durable, local or global in scope;
- a larger knowledge search space than can be explored by any one agent and encouraging economies in acquisition costs
- a variable value of information due to path-dependency of existing knowledge, complicating assessment of the value of any particular bit of knowledge.

These situations are similar to those of all human market participation, while imposing constraints that make their simulation more manageable and controllable for the desired purpose. Simultaneously, they provide the potential for exploring a variety of levels of simulation richness.

The Lobster Fishery as an Example Adaptive ABM

The existing lobster agent based model was created as part of a previous project (Wilson et al 2007) and is now the subject of study and analysis. Its design will be used to briefly illustrate the characteristics of it and the other models being developed.

The LCS

The original idea of a Learning Classifier System can be traced back to John Holland (1986). The LCS used in the lobster model was actually a modified version of this, developed by Stewart Wilson (1994, and Butz 2000) and called an XCS. The XCS is a form of AI consisting of a set of rules and a series of processes operating on the rules. Each rule consists of

- a set of conditions representing the environmental state,
- an action to be carried out in that state,
- an expected reward based on the average of previous rewards under similar circumstances,
- an error representing the deviation of the predicted payoff from the actual reward,
- a 'fitness' that is dependent on the average accuracy of the rule in predicting the payoff,
- and additional fields necessary for implementing the LCS algorithms.

The processes that operate on the rules define means of

- selecting an appropriate action based on environmental conditions,
- modifying or creating rules through genetic combination, random mutation, and covering new conditions,
- selecting those rules best meeting fitness requirements while eliminating rules with poor fitness as rules are tested through repetition and experience.

The result is a system that provides a general capability to 'learn' by developing 'fit' rules based on the agent's experience of the environment and the reward received. In many ways, this system follows closely Herbert Simon's (1955) description of a boundedly rational being learning about his or her environment and making the best decisions possible with the limited experience and knowledge available at each decision point.

The XCS used in the original lobster model was a binary version – environmental state conditions were implemented as strings of bits with three possible conditions – true (1), false (0), or "don't care" (indicated by '#'). The last – the "don't care" state – allowed rules to be created in which one or more of the environmental conditions provided were inconsequential to the decision – either a true or false was allowed. Reward in the lobster model was the catch rate – the number of lobsters caught per trap – and could take non-integer values.

So far, the XCS/LCS described is fairly standard, a typical use of one of the early AI systems in a slightly unorthodox application. However, the manner of XCS use within agents to promote learning is where domain knowledge – knowledge of the fisher's habits and decision processes – enters and melds with AI to create an adaptive agent capable of simulating complex human behavior.

The Lobster Fisher's Decision

Lobsters are relatively simple creatures; their behavior can be modeled on the basis of season, water depth, bottom type, and 'exposure' – a measure of the direction of orientation of an area to the prevailing water activity (waves, currents, and so forth). An extensive knowledge of lobster fishers among some of the researchers and interviews with additional fishers revealed that these conditions – all either easily determined by experience or from nautical charts and other information sources – were the primary determinants of the location of their traps when searching for 'good' spots. In addition, real fishers tended to remember few specific locations and did not usually use means such as GPS units to mark exact locations. Instead, they recalled a general area and its characteristics rather than the latitude and longitude. This general location was embodied in the model by providing an overlay of the model grid that divided the modeled area based on the predominant combinations of the above characteristics. The resulting simplified landscape involved 24 zones, eight months of a lobster season, four depth zones, and three bottom types.

Further complications to an already complex XCS were introduced by incorporating the social possibilities into the model. Although fishing is largely a solitary occupation, coffee shops, restaurants, docks, distributors, electronic communications, and chance encounters at sea provide ample opportunity for social interaction. Interviews indicated that bragging, exaggeration, understatement, and, sometimes, accurate information were topics of exchange during these encounters, and that comparisons based on 'interpreted' results of this interchange were a major means of self-assessment and included questions such as:

- How am I doing compared to my own best?
- How am I doing compared to what I think others are doing?

- How am I doing compared to what I think is the best of the others?
- Has my average catch rate changed and, if so, how?
- How have the catch rates of others changed?
- Should I imitate another fisher and put my trap(s) in the same location(s)?

These questions were factors of fishers' behavior and needed to be included in the model (Wilson et al 2007). These kinds of assessment appeared similar to the 'satisficing' of Simon (1955) but fit nicely into the paradigm of a LCS. However, addition of these new states as part of the human environment complicated the decision process further and increased the single-LCS learning problem by orders of magnitude. A solution to this problem was found by developing a 'hierarchy of decision levels', each with its own LCS, but each using a reduced set of states. Different levels allowed either a final decision (an action to take that resulted in fishing) or an intermediate decision (to utilize another layer of the hierarchy with different states) eventually leading to a final decision.

This hierarchical process mirrors that used by humans in which important, but relatively accessible information, is often used to either adopt a quick choice or to decide that more effort is required for a more in-depth analysis and decision. In our implementation, a socially-oriented, top-level LCS (the Strategy LCS) incorporates frequently-changing social criteria and variability in the perception of relative performance. Simultaneously, this level is accessed at each decision and samples the environment frequently.

A second and third level LCS (the Area and Location LCS levels) then incorporate the more stable biophysical aspects of lobster behavior and the environment, and provide the ability to make a more in-depth decision based on the detailed consideration of past experience in regards to the conditions displaying more regularity. This hierarchical decision process also led in an exponential decrease in the search space of each LCS. The result of this design is shown in Figure 1.

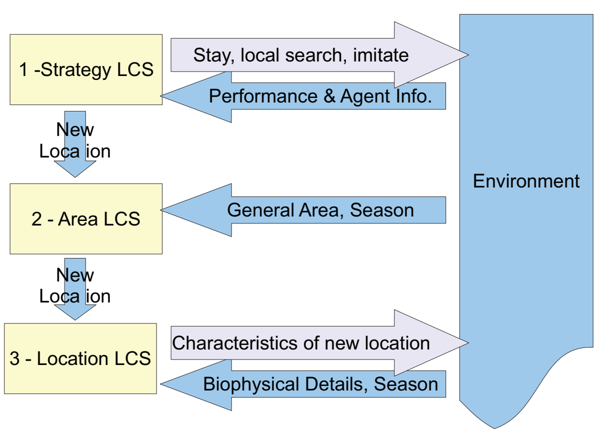

Figure 1: Decision hierarchy of agent.

Social Interaction as a Choice

Lobster fishers (as well as other fishers) interact during their meetings in ways that are often competitive, but sometimes cooperative. Part of our hypothesis is that the positive feedback created by cooperation can lead to sustained social relationships that can, in turn, form the basis for successful management. This appears to have occurred in the Gulf of Maine lobster fishery, which has been operating sustainably for decades under a combination of self and state management (Ostrom 2009). In order to investigate the causes of success and failure, the social interaction needed to be part of the model and incorporated as part of the agent's learned choices.

This was implemented as an action available in the top-level LCS (the Strategy LCS) that allowed one agent to imitate or copy another, laying a trap at the location of the other's recent traps. The decision to imitate is one in which the opportunity cost – either using one's own knowledge or exploring new territory to receive a reward – might be considered in light of the perceived benefit of imitating a particular agent. That perceived benefit would depend on factors influencing the other agent's perceived performance, determined by what was 'heard' about that agent (personal encounters, distributor conversations, the coffee-shop chatter, and so forth) modified by familiarity with that agent obtained through personal encounters.

The binary made representation of such complicated perceptions difficult. The implementation used in the lobster model was to weight the catch rate of another agent by a factor dependent on the relative frequency with which that agent was encountered. This simulated the sense of 'familiarity' developed between pairs of agents. The best weighted catch rate of any other agents was then compared to the agent's own catch rate and the comparison – a binary 'greater than' or 'less than or equal to' – registered as an environmental condition in the Strategy LCS.

Imitation consisted of the imitating agent adopting a location of and setting its trap near one of the imitated fisher's traps. At that point, the biophysical conditions at that location were incorporated as a new imitating agent's rule in the lower two LCS levels, available from then on as any other rule and receiving reward according to performance. In many regards, this is an additional rule creation process of the LCS (although it is largely performed external to the LCS).

Some Preliminary Results of the Lobster Model Analysis

Analyzing the results of the lobster model is an ongoing effort and is leading to additional insight into areas such as the model design, the outputs and information from the model that should be made available, and the limitations of various analysis techniques when dealing with the complex system formed by the model. We present below some of the results in which we have confidence, some suggestions of interesting phenomena, and – particularly in the next section of future work – some of the lessons learned.

Basic Comparisons with Reality

The Maine Department of Marine Resources (DMR) recruited lobster fishers to participate in a program of detailed monitoring in which the location, time, and success of trap hauls was carefully recorded. The resulting database provides nearly a million trap hauls from forty-four lobster fishers over three years and provides the basis for a first comparison of model results with actual performance.

A first measure of performance is the efficiency – the number of trap hauls required to achieve a given production. While absolute value will vary with fisher, location, abundance, and other variables, a telling indicator is the variation of that rate over the months of the fishing season as lobsters perform a seasonal migration and molt new shells. In the real world, this impacts the locations and catchability of lobsters. In the model, this is simulated by a variable catchability factor simulating the availability of lobsters in different habitats at different times of year. Maintaining efficiency in both situations means that fishers must learn to anticipate lobster movement and depletion due to others' fishing. In the model, as in the real world, comparison with nearby competitors leads to a continual jostling of trap locations in search of favorable catch rates as fishers compete in relative performance, resulting in agent adaptation to the lobster environment and a relatively constant haul-per-lobster ratio as illustrated in Figure 2.

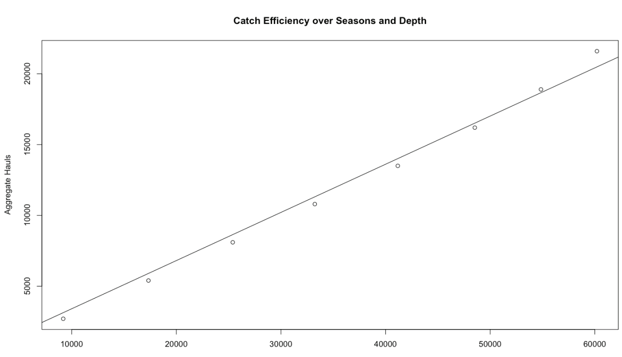

Figure 2: Trap Hauls (vert. axis) vs Lobsters Hauled (hor. axis, demonstrating nearly constant efficiency.

A second comparison made was to plot the lobster catch rate as a function of time over a season; monthly catch rates for the eight-month season are shown in Figure 3. Again, this agrees well with similar plots of real catch rates, and confirms the adaptation of the simulated fishers to the biophysical conditions. In our model, the decreasing catch rate demonstrates a depletion of lobsters as the season progresses.

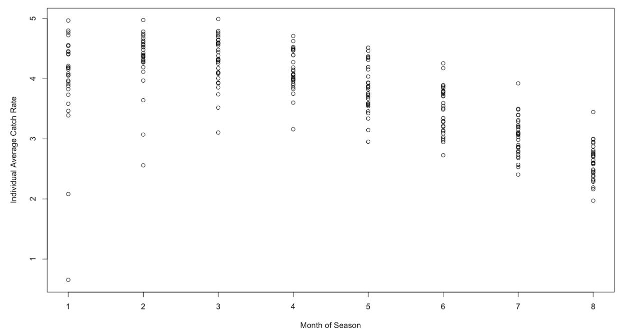

Figure 3: Individual Catch Rates (Lobster/Trap Haul) vs Month of Season, demonstrating a reasonable adaptation to changing lobster habitat.

Social Comparisons

A major component of our interest in creating the model was the potential for studying the formation of social relationships among fishers and the factors affecting those relationships. While we know of no direct or quantitative comparisons with reality that might be used, we know that the social component exists in the real lobster fishery and that the competitive and cooperative actions of participants have served to create, maintain, and enforce a co-management regime with the state that has helped sustain the lobster fishery in Maine.

In keeping with our basic hypotheses that cooperation can help create the basis of this kind of co-management, we have concentrated largely on investigating the networks formed by cooperating individuals who, in the model, form imitating pairs. The model maintains a sociomatrix of imitations that is written to an output text file each year so that, at the end of a typical 50-year model run, there are 50 NxN matrices, showing in the ith row and jth column the cumulative number of times agent i imitates agent j for all model time up to the end of that harvest year.

In general, imitation appears to be a productive tactic for learning about the environment and is relatively popular as a selected action in the top-level LCS. However, the distribution of imitations is not uniform at all. A 'typical' run is shown in Figure 4 from two perspectives. First, all imitations are shown (agents are small circles, arrows indicate the imitations, arrow width is proportional to the number of imitations); next, only those imitations of greater than one per year are shown.

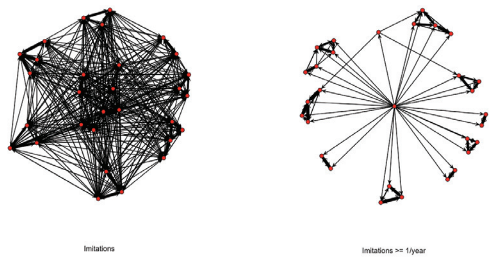

Figure 4: Network diagram showing (left) all imitations and (right) only those > once per year. (Note: placement of vertices in the graphs has no significance, being simply a characteristic of the display algorithm.)

As is obvious, the frequency of imitation varies dramatically among pairs of agents, with a low level (less than one per year) of imitation occurring between almost all pairs, but much more prevalent imitation (up to several thousand imitations over a 50-year run) occurring among some small groups of agents. This is consistent with the types of networks formed by participants in the actual lobster industry where fishers from the same harbor have some closely-knit relationships and other, less-friendly ones, lending credence to the use of this type of model for exploration of human behavior.

On the cautionary side, we have recently determined that the model's implementation of imitation leads to a significant positive feedback and strong path dependence on early conditions in the agents chosen for imitation. We are investigating alternate implementations of the imitated agent selection algorithm in an effort to better understand the present results. Still, despite the strong path dependency that we believe is problematic, there are other tantalizing hints of support for the hypothesis and the ability of these types of models to simulate aspects of human behavior and adaptation, leading us to optimistically pursue continued investigations in this direction.

Ongoing Investigations and Future Work

The Lobster Model

Investigations continue using the lobster model simultaneously with many of the other efforts. Of particular interest is the implementation of the potential for cooperation within the model and the emerging social patterns that result briefly mentioned earlier. However, our post-development work on this model has provided experience in the kinds of information desired and the capabilities that an 'ideal' model developed for these purposes might possess and rewriting has begun on a new iteration alleviating some of the flaws found in the now aging original.

Additional Fisheries: Sea Urchins and Groundfish

Two additional fisheries will be modeled in a similar manner. The sea urchin and the groundfish fisheries in the Gulf of Maine present very different characteristics from those of the lobster fishery while still remaining within the bounds that we might interpret as meeting our overall objectives (as described earlier). Urchins represent an even more sedentary resource than lobsters while groundfish are now recognized as comprising numerous mobile populations with a more limited predictability. Both fisheries are in the process of being developed for modeling using adaptive agents similar to those of the lobster model. The process is similar and consists of

- a round of in-depth interviews with fishers from the industry, researchers, and managers involved in setting policy for the industry;
- transcription of the interviews while distilling aspects significant for simulating the biophysical environment, the social environment, and the decision hierarchy of the fishers;
- discussion regarding the model design and the start of the design, especially the biophysical aspects and the decision hierarchy;
- a second round of interviews with some of the same individuals to better answer questions arising from the previous activities and to propose the basic solution to those in the industry for feedback;
- implementation and tuning of the model;
- analysis and exploratory investigation of results.

The urchin model is the more advanced of the two and has reached the stage of a biophysical model realizing a critical tripping point recognized by both researchers and fishers alike. Urchins, being sedentary, colonize ledges just beneath the low-tide mark (and at lower depths, but with greatly reduced health, nutrition, fertility...). When present, the urchins feed on algae growing on the ledge substrate, keeping algae length short and abundance sparse. When removed from the ledge by a diver during harvesting, algae begin to grow to cover the now barren rock. Urchins (from depths at which they were surviving in poor health or from larvae deposition) may re-colonize the ledge and continue feeding on the growing algae. However, without a nearby population of urchins, and under other conditions leading to high algal growth rates, the algae may grow fast enough to provide shelter for small urchin predators (primarily crabs) that can dramatically reduce urchin abundance or even prevent urchin recolonization. This constitutes a 'flip' or trigger point for the ledge and is a characteristic realized by and to be explored using the urchin biophysical model.

The groundfish model is at an earlier stage. A number of initial interviews have been performed, transcription done, and discussion begun based on the first round of interviews. Although concepts have been developed for the model (both biophysical and fisher) they have not been coded or finalized.

In both cases, work has been a cooperative effort of the interdisciplinary team – including anthropologists, biologists, computer scientists, ecologists, economists, sociologists, and statisticians as primary or temporary members – and results in detailed descriptions of the decision hierarchy to be implemented. Representational details of environmental characteristics, the actions to be provided, rewards to be given, are all discussed at smaller group meetings, and reviewed at project meetings. The belief is that the ultimate product will incorporate the knowledge and expertise of all pertinent domain experts.

A Framework for Future Models

Mentioned briefly earlier, a list of issues, desirable traits and abilities, and some necessary changes were found as a consequence of our investigation of the original lobster model. Primary among these was the recognition that, as the product of many years of development with changing goals, its lifetime was nearing an end. One goal of our continued efforts was to implement a framework to facilitate the creation of new models in the future. After some initial investigation and reviews, we have begun writing wrapper classes around the Mason (Mason 2011) ABM framework to better customize it for our applications.

While keeping the benefits of Mason, our wrappers will provide a very flexible configuration capability from simple text files (for those familiar with it, the ECJ Parameter Database classes are being used and minimally extended), easy creation and provision of I/O capabilities, increased modularity and looser coupling between modules, and greater flexibility in the areas of the environment and provision of 'observer' objects that can fulfill functions common to all agents (such as a common social structure, a marketplace, and so forth). In addition, a configurable integer-based XCS has been added as part of the utilities provided (alternative decision mechanisms may be added as desired).

As with most software, experience will surely prove this to be less than perfect, but we believe it will provide a suite of needed capabilities making implementation and investigation of adaptive agent-based models such as these easier and faster in future efforts.

Acknowledgments

The authors would like to thank the many members of the faculty of the University of Maine and the University of Southern Maine who have provided us assistance in this endeavor. We also wish to thank the many persons contributing to this project in less formal ways, through conversations, suggestions, and, especially, to those who work in the Gulf of Maine fisheries and whose willingness to share both data and their thoughts and strategies have helped make this project possible. Finally, we wish to express our gratitude to the NSF for sponsorship of this effort and to the NOAA NMFS Sea Grant for both their sponsorship of the original modeling work and their continued fellowship support for team members.

References

Arthur, W. et al, 1997 “Asset Pricing Under Endogenous Expectations in an Artificial Stock Market”, Santa Fe Institute Working Paper

Butz,M.V. & Wilson, S.W. 2000 “An algorithmic description of XCS,” lliGAL Report No. 2000017, April, 2000

Holland, J.H. 1995 “Hidden Order: How Adaptation Builds Complexity”, Perseus Books Group, Cambridge, 1995

Holland, J.H. et al, 1986 “Induction: Processes of Inference, Learning, and Discovery”, MIT Press, Cambridge, 1986

Mason, 2011 Multi Agent Simulator of Neighborhoods library, http://cs.gmu.edu/~eclab/projects/mason/

Ostrom, E.; 1990; Governing the Commons: The Evolution of Institutions for Collective Action.; New York: Cambridge University Press.

Ostrom, E.; 2009 “A general framework for analyzing sustainability of social-ecological systems”, Science, vol. 325, pp419-422

Simon, H. “A behavioral model of rational choice”, The Quarterly Journal of Economics, vol. 69, no.1, pp99-118, 1955

Stiglitz, J. 2000. The Contributions of the Economics of Information to Twentieth Century Economics. Quarterly Journal of Economics. pp. 1441-78.

Wilson, J, et al 2007 "The precursors of governance in the Maine lobster fishery". Proceedings of the National Academy of Sciences of the United States of America, September 25, 2007. Vol 104:39

Wilson, s. 1994 “Classifier fitness based on accuracy”, Evolutionary Computation, 1994, vol. 3, no. 2